

~~METHOD AND APPARATUS FOR REDUCING STATIC POWER LOSS~~

Saving Power when in or transitioning to a static mode of a processor

BACKGROUND OF THE INVENTION

Field Of The Invention

This invention relates to computer systems and, more particularly, to apparatus and methods for reducing power use by a computer system during intervals in which processing is stopped.

History Of The Prior Art

As computer processors have increased in ability, the number of transistors utilized has increased almost exponentially. This increase in circuit elements has drastically increased the power requirements of such processors. As the need of power increases, the temperature at which a computer operates increases and the battery life of portable computers decreases. The loss of battery life with modern portable computers greatly reduces the time during which the computer can function as a portable device. In fact, the power usage has become so great that even with significant reduction in the process size utilized, a plethora of techniques have been implemented to reduce power usage to maintain the efficacy of portable computers.

One of these techniques monitors the use of the various devices within the computer and disables those devices that have not been utilized for some period. Because the processor utilizes a significant amount of the power (e.g., 50%) used by a portable computer, this technique is utilized to disable the processor itself when its processing requirements are unused for some interval. In the typical case, disabling the processor is

accomplished by terminating the system clocks furnished to the processor. When processor clocks have been disabled, controlling circuitry (typically a portion of the "Southbridge" circuitry of an X86-processor-based computer) remains enabled to detect interrupts

5 requiring processor operation. The receipt of such an interrupt causes the controlling circuitry to once again enable clocks to the processor so that the processor may take whatever steps are necessary to handle the basis of the interrupt.

10 The technique of disabling the processor reduces significantly the dissipation of power caused by the operation of the processor even at a low frequency. In fact, the technique works quite well; and it is estimated that with many portable computers the processor is placed in the state in which system clocks are disabled during approximately ninety percent of the operation of the computer. However, use of this
15 technique emphasizes another aspect of power loss using advanced processors. When system clocks for a processor are disabled, the processor must remain in a state (sometimes called "deep sleep") in which it is capable of rapidly responding to interrupts. Such a state requires the application of core voltage to the various circuits. The
20 application of this voltage generates a power dissipation referred to in this specification as "static power" usage because the processor is in its static state in which clocks are disabled. To date there has been little attention paid to this static power usage. However, the usage is very
25 significant when a processor functions in the deep sleep mode as much as ninety percent of the time. As process technologies continue to shrink

in dimension and lower operating voltages, this static power increases due to lower threshold voltages and thinner gate oxides.

It is desirable to furnish apparatus and methods for reducing the power use of a processor in the state in which its clocks are disabled.

5 Summary Of The Invention

The present invention is realized by a method for reducing power utilized by a processor including the steps of determining that a processor is transitioning from a computing mode to a mode in which system clock to the processor is disabled, and reducing core voltage to the processor to a value sufficient to maintain state during the mode in which system clock is disabled.

These and other features of the invention will be better understood by reference to the detailed description which follows taken together with the drawings in which like elements are referred to by like designations throughout the several views.

Brief Description Of The Drawings

Figure 1 is a diagram illustrating current-voltage characteristics of CMOS transistor devices utilized in microprocessors.

Figure 2 is another diagram illustrating current-voltage characteristics of CMOS transistor devices utilized in microprocessors.

Figure 3 is a circuit diagram illustrating a first circuit designed in accordance with the present invention for reducing static power usage.

Figure 4 is a circuit diagram illustrating a second circuit designed in accordance with the present invention for reducing static power usage.

Figure 5 is another circuit diagram illustrating a circuit designed in accordance with the present invention for reducing static power usage.

5 Detailed Description

Figure 1 is a first diagram displaying a number of curves illustrating the current-voltage characteristics of CMOS transistor devices utilized in the circuits of a microprocessor. This first diagram utilizes a linear scale for both current and voltage. As may be seen, each of the curves illustrates that the drain-to-source current of a transistor is essentially nonexistent until the voltage at the gate terminal of the transistor is raised to a threshold voltage. Once the threshold voltage of the transistor is reached, drain-to-source current increases either linearly or quadratically depending on whether the transistor is in the linear region or saturation region of operation.

Although the diagram of Figure 1 appears to illustrate that current flowing below the threshold value of the gate voltage is insignificant, this is not the case in some situations. Figure 2 illustrates current versus voltage curves of the typical transistor device below the threshold voltage with the voltage being plotted on a log scale. As may be seen, current in fact flows below the threshold voltage. If a transistor functions in the state below the threshold voltage for ninety percent of computer processor operation, then this current has a significant affect on power usage by the processor.

Since a processor is not capable of computing in the mode in which its clocks are disabled, it would at first glance appear that the solution would be to terminate the application of voltage to the processor.

However, as suggested above, it is necessary that the processor be

5 maintained in a condition in which it can respond rapidly to interrupts provided by the circuitry that controls application of the system clocks.

To do this, the processor must maintain state sufficient to immediately return to an operating condition. Thus, prior art processors have been provided sufficient voltage to maintain such state and to keep their

10 transistors ready to immediately respond to interrupts. In general, this has been accomplished by maintaining the processor core voltage at the same level as the operating voltage. With most prior art processors, the core voltage used by a processor is selected by use of motherboard switches or setup software at a level sufficient to provide the highest

15 frequency operations specified for the particular processor. For example, many processors provide 1.8 volts as a core voltage. On the other hand, the voltage required to maintain state in a deep sleep mode may be significantly less, e.g., one volt or less. Since such processors function at the same voltage whether in a computing or a deep sleep mode, a

20 significant amount of unnecessary power may be expended. In one typical state of the art X86 processor, the power usage averages approximately one-half watt in the deep sleep state because of the leakage illustrated by the diagram of Figure 2.

The present invention reduces the voltage applied to the processor

25 significantly below the lowest voltage normally furnished as a core voltage for the processor during the mode in which system clocks are

disabled thereby reducing the power utilized by the processor in the deep sleep state.

Figure 3 is a circuit diagram illustrating a first embodiment of the invention. In the circuit 10 illustrated, a switching voltage regulator 11 receives an input signal at a terminal 12 which determines its output voltage value. Most modern processors utilize a voltage regulator which is capable of furnishing a range of core voltages for operating transistors; a typical regulator may furnish a range of voltages between 2 and 0.925 volts from which a particular core voltage may be selected for operation. Typically, a binary signal is provided at the terminal 12 which selects the particular output voltage level to be furnished by the regulator 11; in such a case, a number of individual pins may be utilized as the terminal 12.

Recently, a new power saving technique has been utilized which dynamically adjusts both the voltage and operating frequency to a level sufficient to maintain computing operations being conducted by a processor. The technique which offers significant power savings is described in detail in U. S. Patent application Serial No. 09/484,516, filed January 18, 2000, entitled Adaptive Power Control, assigned to the assignee of the present invention. A processor which utilizes this technique monitors operations within the processor to determine the frequency level at which the processor should operate. Depending on the particular operations being carried out by the processor, the value furnished at the terminal 12 of a regulator functioning in such a system will cause the regulator to produce an output voltage at some level

between the high and low values necessary for the particular processor to carry out computing functions.

In the circuit of Figure 3, input to the terminal 12 is furnished via a circuit 13 such as a multiplexor that is capable of providing one or more input values. In the embodiment illustrated, a value is provided at a first input 14 to the circuit 13 by the processor (or other circuitry) which determines the operating condition of the processor in its computing range; and a second value is provided at a second input 15 which is selected especially for the deep sleep condition. Either of these input values may be selected by a control signal provided at a control terminal 16 of the circuit 13. In one embodiment, a system control signal normally utilized to signal entry into the deep sleep condition (a stop clock signal) is used as the control signal to be furnished at the control terminal 16. This control signal selects the input value furnished at the input 15 which is especially chosen to cause a typical prior art regulator 11 to produce a voltage output for operating the processor in the deep sleep mode. In one embodiment of the invention, the value furnished for deep sleep mode is chosen to cause the regulator 11 to produce the lowest voltage possible in its range of output voltages. In one exemplary processor that utilizes the technique described in the above-mentioned patent application, the processor is specified as capable of conducting computing operations in a core voltage range from a low voltage of 1.2 volts to a high voltage of 1.6 volts. On the other hand, the processor when operating in deep sleep mode has no problem maintaining that state necessary to resume computing even though functioning at a core voltage of 0.925 volts, the lowest voltage which the regulator can provide.

Thus, although the voltage regulator 11 may typically provide a range of varying output voltage levels, the lowest voltage at which a processor is specified for conducting computing operations is typically significantly above the lowest value which the regulator is capable of furnishing.

5 In order to reduce power usage in one embodiment of the present invention, in response to a control signal indicating that the processor is about to go into the deep sleep state, the value at the input 15 is furnished by the circuit 13 to the regulator causing the regulator 11 to generate its lowest possible output voltage level for the deep sleep condition. In one exemplary embodiment, the high and low voltages generated in a computing mode are 1.6 volts and 1.2 volts while the deep sleep voltage is 0.925 volts.

10 Although the voltage level furnished by the regulator 11 for the deep sleep mode of the processor might appear to be only slightly lower than that furnished in the lowest operating condition for the exemplary processor, the reduction in power usage is quite significant. Because both the voltage and the leakage current are reduced, the reduction in power is approximately equal to the ratio in voltage levels raised to the power of about three to four. Over any period of processor use involving
15 the deep sleep state, such a reduction is quite large.

20 One problem with this approach to reducing power is that it does not reduce the voltage level as far as might be possible and, thus, does not conserve as much power as could be saved. This approach only reduces the voltage level to the lowest level furnished by the regulator. This
25 voltage is significantly greater than appears to be necessary for a

processor which also dynamically regulates voltage furnished during computing operations to save power. Two criteria control the level to which the core voltage may be reduced in deep sleep. The level must be sufficient to maintain state that the processor requires to function after returning from the deep sleep state. The level must be one that can be reached during the times allowed for transition to and from the deep sleep mode.

The first criterion is met so long as values of state stored are not lost during the deep sleep mode. Tests have shown that a core voltage significantly below one-half volt allows the retention of the memory state of a processor. Thus, using this criterion, it would be desirable to reduce the core voltage to a value such as one-half volt or lower.

However, depending on system configuration, the time allowed to transition to and from deep sleep in an X86 processor can be as low as 50 microseconds. Depending on the capacitive load of the particular circuitry, a voltage variation of about 0.5 to 0.6 volts may take place during this time in one exemplary configuration.

Thus, if the exemplary processor is operating at its lowest processing core voltage of 1.2 volts, its core voltage may be lowered in the time available to 0.6 - 0.7 volts. On the other hand, if the processor is operating at a processing core voltage of 1.5 volts, its core voltage may only be lowered in the time available to 0.9 - 1 volts. Consequently, it is desirable that the core voltage furnished during deep sleep be lowered to a level which may be below the level provided by a typical voltage

regulator but which varies depending on the core operating voltage from which it transitions.

This desirable result may be reached utilizing a circuit such as that described in Figure 4. The circuit of Figure 4 includes a feedback
5 network 41 for controlling the level of voltage at the output of the regulator 11. Prior art regulators such as the Maxim 1711 provide a feedback terminal and describe how that terminal may be utilized with a resistor-voltage-divider network joined between the output terminal and ground to raise the output voltage level.

10 The embodiment of the present invention illustrated in Figure 4 utilizes the same feedback terminal and a similar resistor-voltage-divider network but joins the divider between the output terminal and a source of voltage 42 higher than the normal output voltage of the regulator to
15 force the output voltage level to a lower value rather than a higher level. The particular source voltage and the particular resistor values may be selected to cause the voltage level at the output of the regulator to drop from a particular output value to a desired value such as 0.6 volts when
20 transitioning from a computing level of 1.2 volts.

By appropriate choice of the resistor values of the divider network 41 and
20 the source 42, the voltage drop provided by such a divider network accomplishes the desired result of providing an output voltage for the deep sleep mode of operation that varies from the previous processor computing core voltage by an amount attainable during the transition period available. In one embodiment, resistor 43 was chosen to be 1
25 Kohms, resistor 45 to be 2.7 Kohms, and source 42 to be 3.3 volts. Such

values cause the voltage drop into deep sleep mode to be between 0.5 and 0.6 volts whether beginning at core voltages of 1.2 or 1.6 volts. On the other hand, by using a higher value of voltage at source 45 and adjusting the values of resistors 41 and 43, the increments of voltage drop reached from different starting voltages to final deep sleep voltage values at the terminal 12 may be brought closer to one another.

It should be noted that the circuitry of Figures 3 and 4 may be combined so that both input selection and output adjustment are both used to adjust the core voltage value produced by a voltage regulator for deep sleep mode in particular instances where the load capacitance is relatively low.

Prior art voltage regulators function in at least two different modes of operation. A first mode of operation is often referred to as "low noise" or "continuous" mode. In this mode, the regulator responds as rapidly as possible to each change in voltage thereby maintaining the output voltage at the desired output level as accurately as possible. In order to maintain this mode of rapid response, regulators consume a certain amount of power. When a regulator is supplying a significant amount of power to the load, the power required to operate in continuous mode is relatively small. But, when a regulator is supplying a small amount of power to the load, the power used to operate the regulator in continuous mode becomes significant, and reduces the efficiency of the regulator significantly. It is common for regulators operating in the continuous mode to transfer charge from the supply capacitors back into the power source when the output voltage is changed from a higher voltage to a

lower voltage. The regulator can later transfer that charge back to the regulator output capacitors. Thus, most of the charge is not wasted.

A second mode of operation by voltage regulators is often referred to as "high efficiency," "burst," or "skip" mode. In this mode, a regulator
5 detects the reduction in load requirements (such as that caused by a transition into the deep sleep state) and switches to a mode whereby the regulator corrects the output voltage less frequently. When there is an increase in load requirements, the regulator switches back to the continuous mode of regulation during which more rapid correction
10 occurs. This has the positive effect of reducing the power consumed by the regulator during deep sleep thereby increasing the regulator efficiency and saving system power. But, as a result of reducing the regulator response rate, there is more noise on the regulator output.

It is common for regulators operating in the high efficiency mode to drain
15 the charge on the supply capacitors during a high to low voltage transition on the power supply output or to allow the load to drain the charge. Thus, the charge is wasted during high to low voltage transitions.

It is typical to operate a voltage regulator in the high efficiency mode.

20 Consequently, there is some waste of power whenever a regulated processor goes into the lower voltage deep sleep mode. If the processor is constantly being placed in deep sleep mode, then the loss of power may be quite high. Different operating systems may increase the waste of power by their operations. For example, an operating system that
25 detects changes in operation through a polling process must constantly

bring a processor out of deep sleep to determine whether a change in operating mode should be implemented. For many such systems, such a system causes an inordinate amount of power waste if a processor would otherwise spend long periods in the deep sleep mode. On the other
5 hand, an operating system that remains in deep sleep until an externally-generated interrupt brings it out of that state wastes power through operating the regulator in the high efficiency mode only when the processor is placed in the deep sleep state.

The present invention utilizes the ability of regulators to function in both
10 the high efficiency mode and the continuous mode to substantially reduce power wasted by transitioning between a computing and a lower voltage deep sleep mode. Although regulators have not been dynamically switched between high efficiency and continuous modes, in one embodiment of the invention, an additional controlling input 50 as
15 shown in Figure 5 is added to the regulator for selecting the mode of operation of the regulator based on whether the processor being regulated is transitioning between states. If the regulator receives a control signal 51 indicating that the processor is to be placed into the deep sleep mode, for example, then a regulator operating in the high
20 efficiency mode immediately switches to the continuous mode during the voltage transition. Assuming that the regulator returns the charge to the battery during continuous mode, this has the effect of reducing the waste of power caused during the transition. Once the transition has
25 completed, the regulator switches back to the high efficiency state for operation during the deep sleep mode of the processor.

For regulators that do not conserve capacitive charge by transferring the charge to the battery, a circuit for accomplishing this may be implemented or a capacitor storage arrangement such as a charge pump 53 for storage may be added. Alternatively, when transitioning to deep sleep, the regulator could switch to a mode where the regulator does not actively drive the voltage low but allows the capacitor charge to drain through the load. The selection of power savings modes is dependent on the processor leakage current, the voltage drop between the operating and deep sleep voltages, and the efficiency of the regulator in transferring charge from the capacitors to the power source and then back. If the leakage current is not sufficient to bring the voltage down more than $(1 - \text{efficiency}) * (\text{deep sleep voltage drop})$ during the deep sleep interval, then it is more advantageous to use the load to drain the charge on the capacitors. Otherwise, the charge on the capacitors should be transferred back to the power source.

The control signal utilized may be the same control signal (stop clocks) that signals the transition into the deep sleep state if the method is to be used only for transitions between operating and deep sleep states.

Alternatively, a control signal generated by a particular increment of desired change may be utilized for voltage changes within the computing range of the processor as well as the transition to deep sleep mode.

Although the present invention has been described in terms of a preferred embodiment, it will be appreciated that various modifications and alterations might be made by those skilled in the art without departing from the spirit and scope of the invention. The invention should therefore be measured in terms of the claims which follow.